



HAL
open science

Gaps in Peer Design

Francesca Musiani

► **To cite this version:**

Francesca Musiani. Gaps in Peer Design. Mathieu O'Neil, Christian Pentzold & Sophie Toupin (Eds.), The Handbook of Peer Production, Wiley-Blackwell, pp. 334-346, 2021. hal-03146429

HAL Id: hal-03146429

<https://hal.science/hal-03146429>

Submitted on 19 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

This is the author's version of a chapter accepted for publication in the *Handbook of Peer Production*. Changes resulting from the publishing process such as copy-editing, typesetting, and other quality control mechanisms may not be reflected in this document. This author manuscript version is available for personal, non-commercial and no derivative uses only.

Citation: Musiani, F. (2021). Gaps in peer design. In: M. O'Neil, C. Pentzold & S. Toupin (Eds.), *The Handbook of Peer Production* (pp. 334-346). Malden, MA: Wiley-Blackwell.

ISBN 9781119537106 Available at: <https://www.wiley.com/en-au/The+Handbook+of+Peer+Production-p-9781119537090>

The Handbook of Peer Production

Chapter 25 – Gaps in Peer Design

Francesca Musiani, Centre for Internet and Society, CNRS, France

1. Introduction

In the history of Internet-based services, the concurring pushes towards different types of design choices, in particular peer-based design versus one based on a client/server architecture¹, has been a long-standing source of compromises and tensions – social, technical, political, economic and legal. In a client/server architecture, having a server separate from the user's machine, possibly managed by a third party, greatly complicates resistance to censorship, because it provides an obvious control point (DeNardis, 2014) for authorities. On the other hand, without such a server, communication between users who are not permanently connected becomes much more difficult. This is why fundamental Internet services such as email most often resort to intermediaries, able to ensure the ongoing functioning of the service, but which can also potentially stop, limit, or block this service, as well as read what passes through its servers (on the liability of Internet intermediaries in this respect see Riordan, 2016).

Efforts aimed at developing decentralized systems date back to the early Internet (Minar & Hedlund, 2001). They were generally built as ad-hoc strategic responses to specific threats of shutdown. The filesharing system BitTorrent, for example, was developed as a response to the shutdown of Napster, in order to make legal prosecution for breach of IP in file sharing networks much more complicated (Izal et al., 2004). A system such as Napster, which had a peer-to-peer component in file search, but was in fact being run by a small group of people with associated mandatory passage points for all exchanges, means the system is fully dependent on these nodes, technically as well as

¹ Many Internet protocols are client/server, which means that the machines that communicate are not equivalent: one is a server, permanently on and waiting for connections, the others are clients, who connect when they have something to ask. This is a logical mode of operation, for example, in the case of the Web: when you visit a website, you are a reader, and the entity which manages the website produces the content you seek to read. But not all uses of the Internet fit into this model; they include direct sending of messages, or file exchange – not a one-way communication but a peer-to-peer one, with two machines or two humans communicating directly.

legally. As a consequence, it can be effectively neutralized by turning off servers or seizing the equipment or the people running it (Ku, 2016). In contrast, decentralized systems do not have any central servers, and the functioning of the system involves many peers (people, and the computing resources at their disposal) who do not, or may not, even know each other. If any particular node leaves the system, it continues to run regardless. Thus, the reason for the creation of decentralized systems is that such systems are designed to be resilient to targeting by authorities, and are accordingly often deemed to be superior to proprietary, closed, more centralized systems, because they value long-term robustness over cost-effective commercial expedience (Oram, 2001).

While decentralized systems have been subject to “waves” of interest in the last twenty years, starting with the early 2000s file sharing frenzy and the hailing of P2P as a “disruptive technology” (Oram, 2001), in recent years there has been an even greater interest in, and uptake of, decentralization. In the process, the motivations for adopting decentralized technologies have broadened from a particular strategy of opposition to specific companies or pieces of regulation, to proposals of an alternative “vision” of what corporate, legal and state institutions should be. Two main dynamics have driven this tendency: the first was the emergence of blockchain technology (in particular with Bitcoin technology, as a response to the 2008 financial crisis; see Campbell-Verduyn, 2017 and Brunton, 2019), and the second was spurred by Edward Snowden’s revelations of mass surveillance operations facilitated by a number of telecommunications companies and Silicon Valley giants, on behalf of the U. S. National Security Agency (see Pohle & Van Audenhove, 2017). These events greatly raised the general public’s awareness of the surveillance-based, and personal data-based, business models of near-monopoly tech companies, and their “dangerous liaisons” with state security agencies (Musiani, 2013).

As a consequence of these dynamics, “both decentralization and the notion of authority took on broader meaning and decentralization became a technical, political, economic and social aim in and of itself, reaching outside the ‘hacker’ circles of the early p2p systems” (Brekke & Isakiidis, 2019). However, this wider appreciation of decentralization as a principle and a vision is not devoid of side effects; most notably, oftentimes, decentralization has become an objective in and of itself, with little understanding of intent or assessment of actual effects. As information studies scholar and Internet pioneer Philip Agre said in 2003, “architecture is politics, but should not be understood as a substitute for politics;” decentralized protocols are too readily assumed, because of their technical qualities, to bring about decentralized political, social and economic outcomes.

In this context, it is important to assess not only the qualities, but also the potential problems and difficulties of peer-to-peer and decentralized design, and to assess the extent to which economic, social, legal (and of course, technical) factors may get in the way of the linear “translation” from a peer-to-peer technical architecture to a successful decentralized socio-economic system, able to compete with more proprietary and centralized models. A number of factors may be the cause of this phenomenon, including the difficulty to provide proper quality of service in the early phases of a system that relies exclusively on users’ contributions; the volunteer development model, which oftentimes lacks incentives for performing routine tasks; the difficulty in equaling user-friendly, sometimes non-technical aspects of proprietary systems, such as ‘cool’ design; or the difficulty in finding a straightforward business model that can successfully be merged with decentralization.

The issue of the (de-)centralization of networked structures and organizations has been a central concern for both social and human sciences and computer science when studying the Internet and its numerous social, legal and economic consequences. Beyond

specific discussions of technical architectures and their implications, decentralization connects to broader notions of openness and access in informatics, computing and networking – at the different layers of physical infrastructure, protocols and content – and of information freedom. The emergent field of peer production studies has built upon, and is currently contributing to, discussions around these concepts, drawing from works in anthropology, law, sociology and cultural studies. While this is not the place to review them extensively², it is important to mention their relationship to the present work as they constitute the common foundations on which this chapter, and most chapters in this volume, rest upon.

In the legal studies field, drawing on pioneering works such as those of Yochai Benkler on sharing as a paradigm of economic production in its own right (2004) and of Lawrence Lessig on “code as law” (2002), Niva Elkin-Koren understands architecture as a dynamic parameter and explores how the law does not merely respond to new technologies, but also co-shapes them and may affect their design (Elkin-Koren, e.g. 2006). In the field of economics, Elinor Ostrom (in particular in an edited volume with Charlotte Hess, 2007) demonstrates that collective action and institutional design play as large a part in the shaping of information as do law- and market-based instruments, thereby laying the foundations for new ways of looking at knowledge as a resource shared by a group of people subject to social dilemmas and necessities of choice. From a cultural studies/anthropology perspective, Christopher Kelty has examined the politics of F/OSS by asking about openness the same questions outlined above for decentralization: whether it can be considered as a good in itself, or should be seen as a means to achieve something else, such as interoperability or integration (Kelty, 2008); within F/OSS practices,

² The editors’ introduction to this volume will do so in far more detail; we prefer to focus this chapter on the more specific references related to the “problem” of decentralization and the consequences of peer-based design of the technical architecture.

Gabriella Coleman has examined the role of hacker culture, exploring what hackers mean by freedom and how they enact it, as a form of self-determination that considers unrestricted access to knowledge as a necessary pre-condition for the evolution of their “technical art” (Coleman, 2005). In a more engaged and explicitly activist perspective, practitioners of the F/OSS community, most prominently Richard Stallman, have also outlined the tenets of (networked) information freedom and the consequences of choosing particular sets of technical features, legal licenses, and business models (e.g. Stallman, 2006).

Within this broader scenario, this chapter focuses on the issue of the “gaps” in the peer-based design of the technical architecture of Internet-based services; although net architecture will be our primary focus, we will also consider how dynamics of motivation/incentives to participate in peer-based systems, and their attractiveness/usability, are fundamentally linked to the constraints and opportunities of different architectural designs. The chapter builds upon empirical material derived from two case studies³ of particular events in the development processes of decentralized systems. The first one, the Faroo peer-to-peer search engine, made the core socio-technical choice of placing users and their computing equipment at the core of the system; as a consequence, several strategies were required to remedy technical conundrums and maintain users’ enrollment in the system. The second case, the Wuala distributed file storage platform, made, at a particular point in its history, the choice of reverting to a centralized model, so as to “simplify” the technical development process and the related business model. This revealed a number of technical and social features that pioneer users

³ The fieldwork on which the two cases are based was conducted between 2009 and 2011 and then again in 2014. Results have never been published in English before, but some discussion of them was published in French (Musiani, 2015) and Italian in (Musiani, 2018). The names of developers used throughout the chapter are pseudonyms, although “Faroo” and “Wuala” are not.

were attached to, but that developers were unable to make self-sustainable in order to maintain the peer-to-peer nature of the system.

Via these two case studies, the chapter shows that one of the reasons why these gaps are particularly problematic is that they are not merely ‘technical’ or ‘social’ or ‘economic,’ but a mix of the three. The lessons of these two case studies are also useful to partly explain a widely-known phenomenon related to decentralized architectures, i.e. that while user-controlled, decentralized alternatives to Internet-based services have been regularly put forward as an alternative to internet giants, their developers have mostly found it complicated to compete with proprietary market leaders.

2. Faroo, or the Socio-Economic Implications of a Peer Design Choice

Faroo⁴ was an Anglo-German start-up that, during the late 2000s-early 2010s, proposed a search engine meant to be decentralized at several levels of its technical architecture, based on principles of affinity and common preferences among users. Born in 2007, Faroo was one of the few projects to build a peer-to-peer search engine able to go beyond the stage of thought experiment or university research project. It was led by a small interdisciplinary team including three engineers, a psychologist and a linguist, and assembled a community of about two million users/peers. Its developers, in the early stages of the service, could count on a core of pioneer users, particularly interested in the “full decentralization” model and the “social research”⁵ principles underpinning the search engine. A peculiar characteristic of this service, meant to be P2P, was the choice by its developers to make it proprietary software – which they “justified” by the fact that the specificity of the system was in the technology they developed, not the costs of its

⁴ <http://www.faroo.com>, has now evolved into SeekStorm, a “Search as a Service” project to whose website it is re-directed.

⁵ A number of protocols taking advantage of dynamics of proximity and affinity between users, notably “gossip” and “epidemic” protocols (see https://en.wikipedia.org/wiki/Gossip_protocol).

development, thus, the open source release of the software would be giving the advantage to the competition⁶.

Faroo proposed a “six-times distributed” search, the six levels being indexation, information discovery, search, classification, bootstrapping and software updates, with the aim of building a decentralized network able to free itself from the client-server dichotomy. Pairs were mobilized as components of a socio-technical system entangling users, algorithms and the available technical resources. The fact of not considering peers as merely a support for an intrinsically centralized system⁷, but of “taking seriously” their status of equals within the network and undertaking all development work according to this vision, became a problematic issue in the early development stages of the system. As technical blogger Brad McCarthy commented (2010),

“Peer-to-peer search has some huge advantages. Not being reliant on a single (or even multiple) data center can clearly free up a lot of concern. Socially relevant results do indeed have their place for my own use, and I’m sure they would for others as well. [... At the same time, Faroo] is not being marketed toward anyone. This leaves the entire future of the project (at this moment) in the hands of the user. [...] if you’re relying on the world at large to be your network, and nearly

⁶ The Faroo FAQs, now archived (previously at <http://www.faroo.com/hp/p2p/faq.html#opensource>), read: “Open Source is perfect when competing by a cost advantage with a commercial product on the same technological level (Linux, OpenOffice). But it’s not a good idea to hand over your technological advantage to a monopoly, when competing with its free service with enormous brand power”.

⁷ Some users consider that Skype, in its early phase of peer-to-peer configurations, was actually a system of this kind: “a new node in Skype goes through 2 big steps before it fully joins the network. The first is that it has to authenticate to the central servers, and the second is that it has to find a good ‘supernode’ to initialize/bootstrap and maintain connectivity to the central servers. (...) due to the synchronized restart of clients, when all the nodes started banging on the central servers, the supernodes started losing connectivity and dropped out of the ‘supernode’ mode. Creating a chicken and egg problem and making the situation worse than before. Is it possible that the real reason why the Skype network failed was because of its dependence on central servers?” See comments section of “Skype Failed the Boot Scalability Test: Is P2P Fundamentally Flawed?”, High Scalability, <http://highscalability.com/skype-failed-boot-scalability-test-p2p-fundamentally-flawed>

refusing to entertain the idea of hybrid technology, you may be on the way to kill your company's shot".

Thus, the crucial importance of a peer-centered design took shape at once as the best opportunity and the greatest risk for the Faroo system. Among the opportunities was the fact that the system did not entail significant initial investments of resources by the firm, and was in principle sufficiently flexible to work at once in stable and unstable situations. The system was built to reconfigure autonomously after a potentially disruptive event; in the absence of a central entity, the "surviving" peers were designed to find one another again and ensure, in short delays, the reconfiguration of a distributed entity that could still perform its functions⁸.

At the same time, the centrality of users and of the material and computing resources they would bring to the system could also be problematic for its sustainability, particularly in its initial phases. Faroo "relied on the entire world as a network," as McCarthy (2010) put it. This positioning could be risky, as being able to raise the interest of an important number of users was not only a matter, for Faroo, of reclaiming its niche in the market of search engines, but it was vital for the system to function correctly from a technical standpoint, and for its growth and stability over time: no users meant no resources, in the most fundamental way (Figure 1).

⁸ In a White Paper published in 2007 (formerly available at <http://www.faroo.com/hp/p2p/whitepaper.html>), lead developer Georg used a biology-grounded metaphor to account for peer design: "In biology organisms naturally deal with the rise and falls of its cells; simple elements form superior systems. We believe that evolution works in search too, and that the future belongs to multi-cellular organisms/systems" (p. 22).

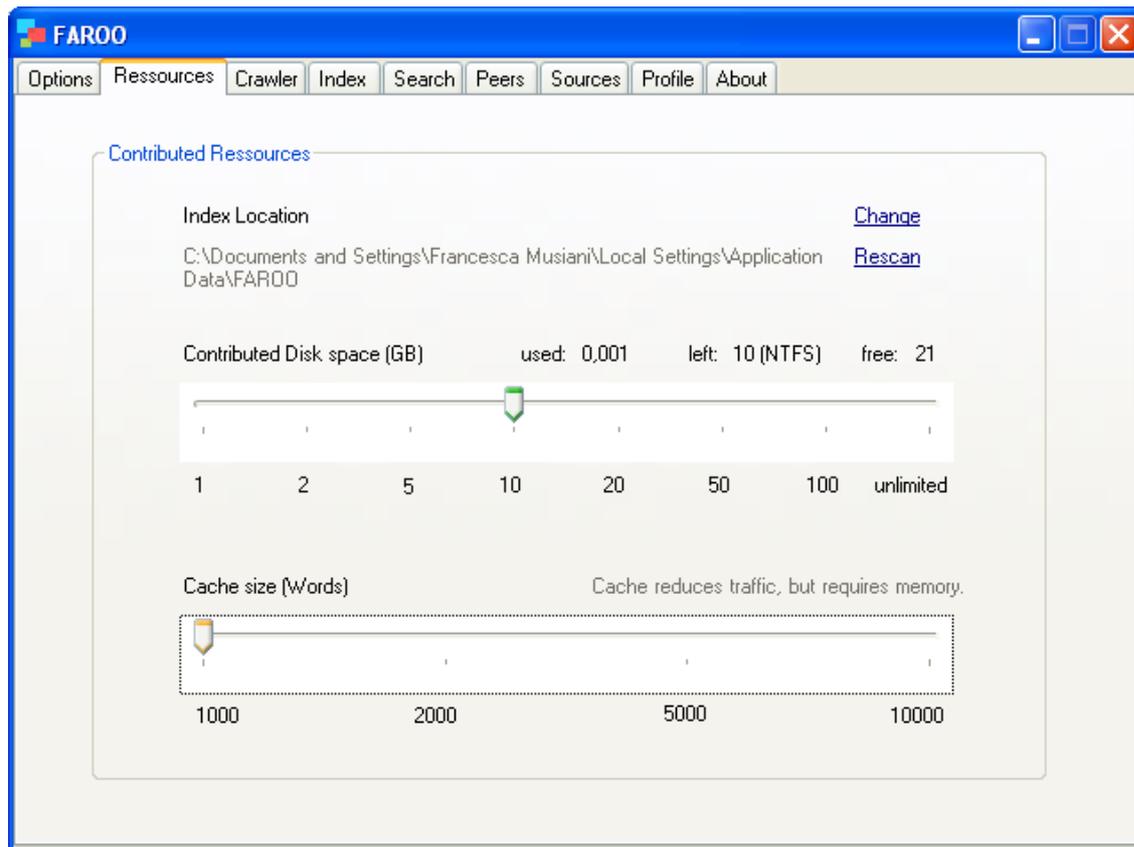


Figure 1. Screen shot of Faroo’s control panel, allowing users to monitor and supervise their contributions to the distributed search engine.

Instead, more users meant not only more material resources (memory, CPU, bandwidth) available to the network as a whole, but also a progressive refining of the variety and trustworthiness of the results obtained by each of the peers, and a better user experience for the network as a whole. This tension was at the core of several points of discussion and controversy during Faroo’s first years of existence.

2.1 “The Egg-Index and the Chicken-Users”: Which comes first?

The relevance and the precision of the search results that could be obtained during the bootstrapping phase of the system – thus, their “quality,” in the view of Faroo’s developers – was the most debated point among the team members: the quantity, and, as a direct consequence, the quality of the P2P search index depended on the quantity of

people/peers using Faroo. During the early beta phase, in 2007, pioneer users considered the quality of the index to be the core problem – to the point of compromising, in some cases, user experience:

“the P2P search index scales with the number of people using FAROO. It looks like a show-stopper. As the number of searchers is very low right now, people are likely to try it, get lousy results and forget about it” (McCarthy, 2010).

The fact that the indexation of web sites was limited to pages that had actually been visited previously by one or more peers was a way to limit spam and improve the quality of the results – but with a low number of “starting” users, the few selected results oftentimes led to disappointing results or even an absence of results. Said for example User 1:

“When you go to perform a search, Faroo will only search sites that real people have visited. This is both the service’s strength and its weakness. You don’t get a lot of junk sites this way. But since Faroo is in private beta with a limited user pool, you don’t get many sites at all that way. A major letdown”.

User 2 summed up this dilemma as the evergreen “chicken and egg problem”:

“As with any social/user-generated software, there’s a chicken and egg problem. It’s hard to market a product that has a small user base, because it just doesn’t perform very well yet. At the same time, it’s hard to attract new users with a product that doesn’t deliver right away all its promises.”⁹

⁹ The highlights of user feedback in this paragraph are issued of a discussion during Faroo’s private beta

ReviewSaurus, a web site displaying technical evaluations of emerging Internet services, proposed a possible solution, highlighting however that it should be temporary if the innovative potential of the P2P index had to be maintained:

“Mahalo¹⁰ displays search results from different search engines. Until they don’t have many users logged in at one time, why not?¹¹”

Facing the possibility that a majority of pioneer users rapidly lose all interest for the project due to lack of proper incentives and motivation to engage in active participation, the Faroo team revised its immediate priorities. The main objective in the phase of “recruitment” of a critical mass of users had to be the scalability of the index, for a rapidly visible improvement of the quality of results, which would have triggered a virtuous circle of motivation and improved technical performances. Temporarily, the technical solution that could be implemented most rapidly was to aggregate, and show to the user, results obtained by other search engines in addition to the results obtained by Faroo. Lead developer Georg underlined that

“(such a solution) is not by itself a big deal and used by other alternative engines, and I wasn’t thrilled. But we had to get stuff that worked immediately, if we wanted to climb on the shoulders of our users, so to speak – and, by getting some results, make the next results better. So, we made a tool with which at some point you

phase, conducted on SeaDVD, a former website where beta-phase software was evaluated (<http://beta.seadvd.com/faroo-follow-up-p2p-search-engine-will-only-work-with-more-users/>).

¹⁰ A now-defunct web directory and human search engine, formerly at <http://www.mahalo.com/>

¹¹ ReviewSaurus, “Faroo: P2P Web search engine, impressive idea!”, January 2008, formerly at <http://www.reviewsaurus.com/web-applications-reviews/faroo-p2p-based-search-engine-impressive-idea/>

could get some results no matter what you were looking for, and you were more willing to continue to use the software. In the meantime, our index was scaling faster.”

The temporary solution produced some beneficial effects for the motivation of the users, and their “recruiting” improved; in March 2008, technical commentators (e.g. on SeaDVD) concluded that Faroo developers have indeed implemented remedial measures, or even that

“they took our suggestion and have started to display the search results from Yahoo or Microsoft Live, Wikia and few others if at all Faroo can’t deliver the search result on its own¹²”.

The design choice led to other “side effects” when it came to the motivation of the users: some of them appreciated that the team of developers was in close connection with pioneer users (“This also shows that they listen to their customers and are ready to customize the application as their user wants”), while other invited explicitly potential users, in forums or in the comment section of specialized blogs, to join the system.

Interestingly, an argument supporting this “call for users” was the fact that a peer could join with different degrees of engagement, at different times, all the while helping the system to bootstrap. Indeed, from the very start, and even as a passive user (i.e., installing the Faroo P2P client even without actively launching search requests), a peer already contributed to bring resources to a system that was in great need of “globalizing” itself:

¹² ReviewSaurus, “Faroo opens for public and takes our suggestion!”, March 2008, formerly at <http://www.reviewsaurus.com/software-reviews/faroo-opens-for-public-and-takes-our-suggestion/>. The following two citations are also from this webpage.

“a suggestion for all of you too, go ahead and install Faroo and use it as much as you can, it’s high time and we need an alternative of Google and you can help in doing that by just running this client.”

Thus, it was possible – if the user was unsatisfied by the performance of the search engine once the client was installed – to let it run in the backstage for a few months and then come back to it, when more peers had joined the system and as a consequence, the index had grown in size and complexity. According to the technical blog ReviewSaurus, referenced above, this was “worth a try,” especially as the quantity of resources mobilized by Faroo did not reduce the reactivity of the user’s terminal, a problem that is often caused by software functioning “locally,” such as P2P clients. This feature became, in the eyes of the users, a gesture of “technical politeness” which prompted them to recruit other users. SeaDVD’s review of Faroo remarked that

“the client doesn’t seem to eat much RAM, so you could always install it now, go about your business and check out the search engine feature in a few weeks or months”.

While the episode shows that some strategies can be implemented to temporarily remedy it, the “chicken and egg” problem of users and their resources is indeed an important hurdle that peer-to-peer based design has to face.

2.2 Downloading a P2P client: A barrier to entry?

A second point of controversy in the Faroo case which is illustrative of a problematic aspect in peer-to-peer design concerns the status of the P2P client¹³, that Faroo users needed to install as a precondition to becoming a part of the system. In Faroo's early development phases, the P2P client was for the developers a central artifact, both materially and symbolically: the concrete "place" and technical device where the core operations of indexation and search happened, and the symbol of the "distributed centrality" of users in the system. However, at the moment when the closed beta version (a relatively stabilized version, accessible by developer invitation) of Faroo was released in early 2008, users issued frequent objections to the use of client software which was to be installed on their terminals, citing two recent (at the time) tendencies in the market of Internet services: the rise of cloud computing, and an increasing number of users worldwide shifting to exclusively mobile Internet connections. Said a user on a Faroo-related discussion thread:

"Faroo, being P2P, requires a client download. That can be an incredibly big hurdle to adoption. At first I thought your best shot was to build on top of existing P2P services such as LimeWire or Gnutella. But your strategy appears to be different as you are building on top of .Net.... Oops there go all the Mac and Linux heads... bye. This is potentially a serious issue as lots of early adopters tend to be on the Mac. If I am a Mac user, then I may not use Faroo without borrowing somebody's PC"¹⁴.

¹³ So as to clarify all semantic ambiguity, it should be pointed out that throughout this section the word "client" is used to indicate the piece of software that a user needs to install on her machine so that it can become a node in the P2P system.

¹⁴ User feedback in the commentary section of Lunn, B. "Faroo: Could P2P Search Change the Game?", ReadWriteWeb (now ReadWrite), Web technology blog formerly at http://www.readwriteweb.com/archives/faroo_could_p2p_search_change_the_game.php. The following two

Furthermore, the fact that the P2P client was developed on NET (an app-building kit of tools and libraries developed by Microsoft, thus, strictly linked to the Windows environment) seemed to be a deal-breaker for some pioneer users. The most concise one, Fabian, summed up “No Mac, No Linux? Then not for me”. Another user, Samidh, in a less strict but equally critical fashion, remarked that the multiplication of devices linked to the Internet that are now owned by the average user, and the current tendency to make terminal-installed applications increasingly “light” and mobile, could cause supplementary problems for the diffusion of an application that implies a mandatory client download:

“I think the biggest obstacle is that they require a download. Startups need to realize that the age of the download is dead. Skype was the end of an era. People are far too mobile and use too many computers to worry about specific downloads anymore. Everybody now needs to keep repeating the mantra, ‘The Internet is the Platform.’”

And danah boyd, in her then-capacity as a Microsoft researcher¹⁵, contributed to present the obligation to install a P2P client as a “weight” on the shoulder of potential users, rather than emphasizing potential advantages such as resilience:

“I’m also curious about the future of peer-to-peer systems in light of the move towards the cloud, but I’m not convinced that decentralization is a panacea to all of our contemporary woes. Realistically, I don’t think that most users around the globe

interventions are also hosted on this thread.

¹⁵ Thus, it should be noted, not a completely neutral position in the centralization/decentralization debate....

will find a peer-to-peer solution worth the hassle. The cost/benefit analysis isn't in their favor.”¹⁶

Thus, the P2P client, symbol of the articulation of the global and the local dimensions of a P2P system, and “sold” by the Faroo team as one of the innovative and original features of the system, became, through the feedback of several users and technical commentators, an obstacle and a barrier to entry – something that prevented an immediate and comfortable appropriation of the system by the users. The Faroo developers were aware of the issue: one of them, Jiechi, placed the difficulties the service faced within a broader landscape of similar problems faced by other P2P systems:

“Indeed, it is not easy even for successful ones, look at Joost¹⁷ – they've been forced to move to a pure browser model because nobody took the download plunge. The only people who can successfully get downloads distributed are people who already have wide distribution.”

However, the P2P client download was, in lead developer Georg's vision, one of the distinctive, and “deal-breaking” features of a decentralized search engine, being the way in which users could monitor the incoming and outgoing traffic of its P2P client – traffic which was an indicator of two things: that the search query was being indexed by the peer that the client represents, and that this peer was in the process of answering queries. Thus, he engaged in a “translation” attempt, and redefined the P2P client as a benefit for the user:

¹⁶ danah boyd's contribution on the P2P Foundation mailing list, May 27, 2010, formerly at http://listcultures.org/pipermail/p2presearch_listcultures.org/2010-June/008941.html.

¹⁷ Joost was an Internet TV service created by Niklas Zennström and Janus Friis (previous founders of Skype and Kazaa), which during a period of its development used peer-to-peer television technology for content distribution.

“With the bandwidth rise downloads are not that uncomfortable anymore. [...] ‘The Internet is the Platform’. So do we [try to do], referring to the decentralized structure of the internet, including users as part of the internet, not only reducing it to a bunch of central servers ;-) I believe that people will jump over this ‘download hurdle’ if they get something in return, which they can’t get in a centralized system...”

This point of view was shared and clearly communicated by Bernard Lunn, a technical commentator who would later become a Faroo consultant:

“The download hurdle is a common reservation. It has not proved a hurdle when the payoff is big enough – think Skype and Spotify. And more recently the iPhone and Android app stores have totally changed the mindset around downloading software. The key issue I think is what you mention – “hassle”. It is a hassle-free experience on iPhone but people have found the experience less than hassle-free on the PC.¹⁸”

Thus, from a narrative of gaps, hurdles and problems, another one sought to emerge, describing the download of a client that could be hosted on one’s terminal, assuring the P2P connections, as an economic opportunity and a market niche, the symbol of another way to conceive information search on the Web and maybe, eventually, the Internet at large. Indeed, at stake was the interoperability of P2P applications with the rest of the network:

¹⁸ Bernard Lunn’s contribution on the P2P Foundation mailing list, May 27, 2010, as an answer to danah boyd’s previously examined remarks.

“The first P2P application was file-sharing, the next Skype, now we are building P2P search. What could be the next? Instead of another isolated P2P application, we would like to see P2P built into the OS and Internet stack in a standardized manner.¹⁹”

In the end, Georg summed up, the download and local running of the P2P client could be framed as either a gap/problem, or an opportunity. What would tip the scales on one side or the other was the added value that the decentralized system would be able to provide, both in terms of business model and user experience: “... this will be privacy protected personalization, attention based ranking, and revenue sharing based on the saved infrastructure costs”. But once again, the controversy around Faroo’s client download demonstrates the technical and socio-economic costs of making the peer design choice, and the need for Faroo’s developers to engage in a number of strategies in order to keep the product attractive for users with respect to centralized platforms.

3. Wuala, or Ultimately Renouncing Peer Design

Wuala²⁰ (2007-2015) was a start-up born in Switzerland, later acquired by a leading French manufacturer of storage hardware, offering a distributed file storage system that also included social networking features. Created in a late-2000s landscape replete with online storage applications based on “classic” servers or cloud platforms, Wuala’s team of developers - fifteen at the time of my initial fieldwork in 2010, all computer engineers – worked for a number of years on a hybrid storage platform, combining centralization and

¹⁹ For this and the following citation, see Faroo White Paper at p. 30.

²⁰ Formerly at <http://www.wuala.com/>

decentralization, and advocating a sophisticated system of “bartering” disk space in exchange for permanent storage space on the user network. Wuala was an interesting example of a “compromise” between an approach largely dependent on a server platform, and a distributed and decentralized architecture approach. It proposed a decentralized version of an online storage platform, of which Dropbox currently is the leading centralized representative. Furthermore, it possessed a relatively large user base actively engaged, like in Faroo’s case, in the discussion and shaping of the decentralized aspects of the system.

I analyze here a specific moment of “gap” in the P2P-based design of Wuala, the time at which – on September 28th, 2011 – Wuala developers chose to interrupt one of their defining features, the “barter” of disk space described above. This decision *de facto* meant the “re-centralization” of the system at the level of the technological architecture, depriving it of its fundamental distributed feature – the feature that made the user’s computer an active part of the archive system, rather than the mere “starting point” for files which end up in central servers. The decision also meant ending what was universally regarded as the main specificity of Wuala compared to similar services available on the market. The developers were not insensitive to the controversy that this might cause, and tried to prevent it through the following post on their main user forum²¹, emphasizing the positive implications of the choice not only from their point of view, but what could be the advantages for users:

Why is trading not supported anymore?

²¹ Wuala’s user forum was hosted on the start-up’s website, www.Wuala.com, and is no longer accessible since November 15th, 2015, when the service was shut down.

With the introduction of our new storage system, the whole file storage scheme has changed. There are some major improvements coming with this new storage system, such as incremental upload, higher redundancy, multiple down- and uploads and increased file size limit (up to 100 GB). Also, we want to free up engineering resources so we can focus on features we consider more important. (Gianluca, Wuala developer, September 28th, 11:53)

The discussion that follows is an exhaustive testimony of how the re-centralization, and to some extent the failure, of the peer-to-peer project in Wuala – which seemed to betray the system’s original intent and “reform it” according to the dominant model, so much so that it was for many a deal-breaker – not only responded to a variety of motivations from the developers, but was understood in different ways by the user community, and affected their perception of whatever incentives were left to them for participating in the system. On the basis of the issue of re-centralization, the discussion about a specific system and its evolutions became a broader one about the pros and cons of peer-based design – from ‘practical’ and technical aspects to issues of principles, of community, of organization of the Net.

3.1 The many facets of discontinuing a peer-based feature

First of all, the end of the peer-based feature was framed in terms of the loyalty that developers owed to users in exchange for their active contribution to product improvement. In fact, as always happens in the initial phases of the development of a totally or partially open source software, the first Wuala users and their computers served as a ‘test-base’ for the system, highlighting, through their practice of it, a series of aspects that could be improved. The re-centralization therefore signified the failure of the trust that users placed in the developers – trust that they were contributing to something of value, different from

most of the solutions proposed on the market, and whose creation process really “took seriously” the community on which, socially and technically, the system was based:

As a user from day one it makes me really sad and angry to see this change.

During the last months Wuala has already stopped listening to its loyal users and with every update there were missing features, incomprehensible design decisions, etc.

And now - booom! [...] You tricked the community in believing they have a saying in your decisions. I want the time and effort I spent with testing, commenting, suggesting and bug tracking over the last years back from you!! (joe12, October 3rd, 5:15pm)

My guess is that a lot of users will feel kind of mis-used [...] And as fast as new products and services can appear in the internet they can also disappear. I would think that the impact on wualas image is much bigger than it looks right know: one key essential when using cloud-based services is trust. Changing the game in a way that it looks solely money-driven will disappoint the community that also supported wuala a lot in the past. (Doug0915, October 3rd, 5:55pm)

A related issue, but with its specificity, concerned the possible economic models for peer-to-peer. Various users seemed to suggest that the barter system was a victim of the impossibility of continuing to maintain a peer-based model when it comes to profits. Was Wuala’s story that of many originally peer-based systems, that were born peer-to-peer to grab some “loyal” users seduced by the model, and that ended up succumbing to the temptation of improving its economic performance by introducing a simpler model

technically, but which made it easier to enter mechanisms that could be sources of income for the company, such as advertising or premium solutions?

No trading = no wuala for me. Thanks for the run -- this was a great idea that seems to have been ruined by a corporation. I understand that everyone is just in this game to make money, but it is sad to see a truly useful feature disappear from this product. (VenMen, October 3rd, 5:09pm)

I understand you want the money, but imo this change will reduce long-run income (no more advertising from guys like me who is/was in love with the system). Of course it will give more money in a short-term. Probably that's your goal, maybe you are not planning to live long. (Dimps, October 3rd, 5:48pm)

One of the fundamental aspects that seemed to fail with the re-centralization was the attribution of independence, autonomy and control to the users allowed by P2P design, with the strong “local” dimension of the storage and the sharing of each user’s disk space with other users. Many had chosen Wuala precisely for this distributed character, which allowed to free oneself as much as possible from the service provider and therefore to be as independent as possible from servers and the terms of use they imposed, by replacing them with a network of peers. A user went even farther, explicitly mentioning the failure of the company’s “social experiment,” which therefore invalidated in the eyes of users any solution that could be proposed later.

[The reason] that I selected Wuala is gone for me. The whole idea of having data stored on several peers (which improves speed etc), and sharing out one's own disk

accordingly, was what appealed to me and what I had been looking for all since this idea originated. [...] The very strong argument of having the data distributed across many computers in the net made Wuala so appealing, giving the impression of independence and not rely on a single company for storing all the data fragments. (flexgrin, October 3rd, 6:12pm)

[Now you could at least] release a Wuala based client [...] using the old scheme as open-source, so that new clouds of **social** storage can emerge using it. Not doing so basically declares the attempt to implement a social file storage system as a failure, and if it was a failure, why trust the company behind it? (Tor, October 3rd, 5:23pm)

And for others still, renouncing peer design was a failure because the peer-to-peer model embodied various forms of “technological activism” – environmental protection (using free and underutilized disk space on their computers instead of feeding the data centers with content duplicates), being part of a community interested in the social value of sharing hardware and software resources, or even referring to the original (and idealistic) principles of the Net such as cooperation, symmetry, equality:

Maybe I have been operating under a happy delusion for the past couple of years, but I thought that by sharing my storage with other Wuala users, I was, in some small way, contributing to the health of our planet. I thought I was helping to reduce power consumption by making it possible for your company to store data on underutilized storage space on my PC and notebook computers. (droom, October 3rd, 6:09pm)

There was always this ‘promise’ that one can be part of the storage community by sharing space. Many wondered whether this could sustain but the answers I remember have always been confirming the commitment to sharing space as being one of the core principles of Wuala. I am disappointed because I thought the trading option was a very nice social experiment marred by a flawed implementation. I am a strong believer in distributed systems at the endpoints of the net. (Milic, October 3rd, 6:45pm)

4. Conclusions

This chapter has addressed the issue of the “gaps” in peer-based design of the technical architecture of Internet-based services, by building on two case studies of particular events in the development processes of decentralized systems. Via these two case studies, the chapter has sought to address the “mix” of technical, social, and economic factors that lead to dilemmas and problematic issues in peer design.

With the case of Faroo, we saw how the fact of considering peers not as mere outposts of a centralized system, but taking “seriously” their status of equals at the network’s periphery, and having this design as foundation of all further work, represented an extremely delicate issue for developers at the early stages of system’s expansion. We also saw how the conception of the user’s “centrality” (and the resources that it brings to the whole system) resulting from this distribution approach emerged as both the greatest opportunity and the major risk for the system. Among the advantages, such a system did not require large investments of resources from the company, and could work in stable and unstable situations. It was also possible for it to recover autonomously after an event having changed its structure, thanks to the “surviving” peers that guaranteed the readjustment of a distributed entity that could still perform its functions. At the same time,

the centrality of users and their computing resources could also be problematic for the sustainability of the system, particularly in its initial phases. To rely on no other resources than those provided by users, “to rely on the world to be its network”, also meant, for developers, to take a huge risk, betting on the fact that they would manage to “enroll” (Callon, 1984) a number of users large enough to ensure the proper technical operation of the system itself, and to allow its growth and stability over time. More users meant not only more material resources (memory, CPU, bandwidth) available to the whole, but also, as we have seen, a progressive refinement of the variety and reliability of the results obtained by each of them, and a better user experience for the entire network. We have followed how this tension took shape around two points of controversy within the team, and between the team and some users, during the first years of Faroo’s existence.

The first of these cases, concerning the quality of the results during the bootstrapping phase, showed how – given that the quantity and quality of the P2P search index depended on the amount of people/peers using Faroo – the shortage of users proved to be the main problem, both human and technical, for early adopters, a situation which forced developers to temporary technical fallback solutions. The second case followed discussions focused on the P2P client, which Faroo users had to install as a precondition to entry into the system. The symbol of the articulation between the local and global dimensions of the system, understood and “sold” by the Faroo team as one of the innovative and original points of their service, was transformed, by the reactions of some users and commentators, into an obstacle, a barrier to entry, forcing the Faroo team to renegotiate the P2P client as a market opportunity, symbol of another way of conceiving the search for information on the Web and perhaps the Internet as a whole.

In the case of Wuala, we have witnessed a moment in the history of a P2P system in which its developers decided to renounce decentralization, and discontinue peer design.

We have seen how changes in the design of the technological architecture of a communication network has both causes and consequences that are economic, political and social; when networking models, the management of data flows, their archive and technical treatment are discussed, the peers no-longer-to-be were in fact discussing of these broader consequences, their perception as more or less just and fair, socially and economically, and more or less successful. The ability of the network (and its developers) to sustain peer design became the core notion around which dynamics of equality, of “justice,” of community took shape. The re-centralization of the system was the focus of the negotiations between innovators and pioneer users, and it showed the reconfiguration and re-appropriation of peer-based innovation, creating multiple definitions of what constitutes the technical, economic, and political “failure” of a P2P system that is not such anymore.

This process also shows that the “barter” feature in the Wuala system had, by the time it was discontinued, assumed the quality of “distributed governance,” implemented through peer-to-peer technology. When the feature was suspended, users felt deprived of this management model by an act of politico-economic governance, giving predominance to a simpler business model, which is perceived as more dictatorial and oblivious to the wishes expressed by users, as well as the care they have displayed to make the previous system a better technical artifact, driven by practical and principled motivations. The choice to relent peer design becomes a symbol of this evolution of power and control relations, as well as trust. The evaluation of the soundness of the different architectural solutions – the previous one and the new proposal – with respect to Wuala’s business model, which the developers thought they could make exclusively based on technical and

economic considerations, became charged with other meaning and eventually caused the system's demise²².

In conclusion, assessing the advantages and disadvantages of peer design is a foremost example of a long-standing heavily discussed issue: whether a technical architecture should be centralized or decentralized, its design based on equal interaction among peers or on a client/server architecture. It is important to underline that, when it comes to discussing the degree of (de-)centralization, protocols are not intrinsically good or bad. The two cases we have examined – the gaps and shortcomings they show when it comes to peer design, and the strategies that developers and users display to counter them – ultimately illustrate that choices related to the design of technical architectures “come with” a number of social, organizational and at-large political dynamics, and lead to stories of compromises and trials²³.

However, they also reinforce Philip Agre's seminal argument that while architectures are politics, they should not be assumed to be a substitute of politics (Agre, 2003). When it comes to privacy, user autonomy, and freedom of expression, Silicon Valley giants and their centralized silos may arguably not be the best option, but having a system with intermediaries and obligatory passage points may actually prove more useful, in some cases, than a completely egalitarian and peer-based design, and in particular, it can be understood by developers-turned-start-up-leaders as a simpler way to go for elaborating a business model. And despite an accompanying rhetoric of openness and freedom that present them as the ideal solution to all problems, including censorship, decentralized architectures also bear or can bear issues.

²² On August 17th, 2015, Wuala announced that it was discontinuing its service and that all stored data would be deleted on November 15th, 2015.

²³ In the sense of *épreuve* (see e.g. Martuccelli, 2015).

This is not only because of socio-technical and socio-economic constraints, but also – being potentially subject to issues of data poisoning, net neutrality, display of metadata such as IP addresses, and difficulties to enforce identity verification – from the standpoint of civil liberties and human rights. While the present chapter has focused more on the former aspects, and what they mean for the capability of peer-based systems to compete with centralized ones, we should not neglect the latter as they are, as well, an integral part of how Internet governance is increasingly conducted via its infrastructure. Together, all these implications of architectural design choices inform what we have called “the turn to infrastructure in Internet governance” (Musiani et al., 2016), and as a consequence, the kind of Internet user that we can and want to be.

References

Agre, P. (2003). Peer-to-Peer and the Promise of Internet Equality. *Communications of the ACM*, 46(2), 39-42.

Benkler, Y. (2004). Sharing Nicely: On Shareable Goods and the Emergence of Sharing as a Modality of Economic Production. *The Yale Law Journal*, 114 (2), 273-358.

Brekke, J. K. & Isaakidis, M. (2019). Principles for Decentralized Protocol Design. NEXTLEAP Project Deliverable 2.5.

Brunton, F. (2019). *Digital Cash: The Unknown History of the Anarchists, Utopians, and Technologists Who Created Cryptocurrency*. Princeton University Press.

Callon, M. (1984). Some elements of a sociology of translation: domestication of the scallops and the fishermen of St Brieuc Bay. *The Sociological Review*, 32(1-suppl), 196-233.

Campbell-Verduyn, M. (Ed.). (2017). *Bitcoin and Beyond: Cryptocurrencies, Blockchains, and Global Governance*. Routledge.

Coleman, E. G. (2005). *The social construction of freedom in free and open source software: Hackers, ethics, and the liberal tradition*, PhD Thesis, Chicago, IL: The University of Chicago.

DeNardis, L. (2014). *The global war for internet governance*. Yale University Press.

Elkin-Koren, N. (2006). Making Technology Visible: Liability of Internet Service Providers for Peer-to-Peer Traffic. *New York University Journal of Legislation & Public Policy*, 9 (15), 15-76.

Izal, M., Urvoy-Keller, G., Biersack, E. W., Felber, P. A., Al Hamra, A., & Garces-Erice, L. (2004). Dissecting bittorrent: Five months in a torrent's lifetime. In *International Workshop on Passive and Active Network Measurement* (pp. 1-11). Springer, Berlin, Heidelberg.

Kelty, C. M. (2008). *Two bits: The cultural significance of free software*. Duke University Press.

Ku, R. S. R. (2017). The creative destruction of copyright: Napster and the new economics of digital technology. In *Copyright Law* (pp. 207-268). London: Routledge.

Lessig, L. (2002). *The Future of Ideas*. New York: Vintage Books.

Martuccelli, D. (2015). Les deux voies de la notion d'épreuve en sociologie. *Sociologie*, 6(1), 43-60.

Minar, N. and Hedlund, M. (2001). A network of peers – Peer-to-peer models through the history of the Internet. In A. Oram (Ed.), *Peer-to-peer: Harnessing the Power of Disruptive Technologies*, 9-20. Sebastopol, CA: O'Reilly.

Musiani, F. (2013). Dangerous liaisons? Governments, companies and Internet governance. *Internet Policy Review*, 2(1).

<https://policyreview.info/articles/analysis/dangerous-liaisons-governments-companies-and-internet-governance>

Musiani, F. (2015). *Nains sans géants: Architecture décentralisée et services Internet*. Paris: Presses des Mines.

Musiani, F. (2018). “Quando il peer-to-peer si ‘ricentralizza’: Vincolo sociotecnico, spinta di mercato o fallimento?” In Balbi, G. & Magaudda, P. (eds). *Fallimenti digitali: Un’archeologia dei nuovi media*. Bologna: Unicopli.

Musiani, F., Cogburn, D. L., DeNardis, L., & Levinson, N. S. (eds.). (2016). *The Turn to Infrastructure in Internet Governance*. Springer.

McCarthy, B. (2010). “A look at Faroo’s P2P Search. Will this one make it?”, *The Next Web*, May 25. <http://thenextweb.com/apps/2010/05/25/p2p-search-is-here-but-who-will-use-it/>

Oram, A. (Ed.) (2001). *Peer-to-peer: Harnessing the Power of Disruptive Technologies*. Sebastopol, CA: O’Reilly.

Ostrom, E. and Hess, C. (2007). *Understanding knowledge as a commons: from theory to practice*. Cambridge, MA: MIT Press.

Pohle, J. & Van Audenhove, L. (eds.) (2017). Post-Snowden Internet Policy. *Media and Communication*, 5(1).

<https://www.cogitatiopress.com/mediaandcommunication/issue/view/59>

Riordan, J. (2016). *The liability of internet intermediaries*. Oxford: Oxford University Press.

Stallman, R. (2002). *Free software, free society: Selected essays of Richard M. Stallman*. Lulu.com.